VIEWPOINT May 7, 2007, 12:01AM EST

The Age of the Supercomputer

Ten years ago, Deep Blue beat a world chess champion. Now supercomputers are poised to move into the business world in a big way

by Dave Turek

On May 11, it will have been 10 years since Deep Blue became the first machine to defeat a reigning world champion chess master under tournament rules. The eight-day, six-game match with Garry Kasparov drew the world's attention. It was parodied by late-night comics and in comic strips. It became fodder for folk singers, documentary producers, and book authors. Deep Blue, a machine built by IBM (IBM), was a pop star.

Things have changed in 10 years. Today, Microsoft's (MSFT) Xbox 360 game console is more powerful than Deep Blue, and the real silicon speedsters only make headlines for the fleeting moment when they wear the title, "Fastest Supercomputer in the World."

Deep Blue was not the fastest computer in the world in 1997. On a list of the top 500 fastest computers, Deep Blue ranked a poky 388. Still, its 32 processors could handle about 11 billion calculations per second.

Ten years and four processor generations later, Blue Gene, the present fastest supercomputer and the progeny of Deep Blue, uses 131,000 processors to handle 280 trillion operations every second.

THE BEAUTY OF VIRTUAL MODELS

The biggest advance isn't the technology per se. A decade since Deep Blue, the business world is now witnessing the advent of fully commercial supercomputing. The real revolution is yet to come.

Supercomputing is moving rapidly from its traditional place in academic and government research to tackle a growing list of commercial applications in business. Of the top 500 supercomputers in existence in 1997, only 161 were employed in businesses. Today, nearly half of the 500 speediest—246 systems—reside in commercial enterprises.

Two essential economic forces are propelling that movement. The first is the ever-rising expense of building and testing physical prototype designs for everything from diapers to cars. The second is the rapidly decreasing cost of perfectly modeling and testing a virtual product on a supercomputer.

In 1997 the price for the capability of processing 1 million operations per second was roughly $50. At that rate, Deep Blue's 11 billion calculations cost about $550,000.

Today, 1 million ops cost about 10 cents—or $1,100 for the same computing capacity. As a general rule, supercomputer power increases and prices fall about 40% each year. The fastest supercomputer in the world in 1997, a multimillion-dollar capital expenditure, is today's $200,000 starter kit.

MODELING FINANCIAL RISK

Experienced users of supercomputers are riding the price/performance curve to apply the machines' incredible power to areas that not so long ago would have been considered too esoteric or tangential to waste valuable resources on. Procter & Gamble (PG), for example, has become so adept and efficient in its use of supercomputers that it can afford to tweak the design of Pringles potato chips so they fly more efficiently through the production line; make diapers more eco-friendly and absorbent; and make soap suds bubblier using computer modeling and simulation techniques that until recently were the province of aerospace engineers.

Early pioneers, in industries where massive computer power has long been justified as a necessary capital expense, will continue to demand machines that push the limits of computer theory. As the cost of drilling a single well in deep water approaches $35 million, we can expect the energy industry to further expand the boundaries of computational modeling of petroleum and gas reserves. Some of the new applications that will emerge among large users lie in managing digital media and modeling financial risk.

Many big companies don't know it yet, but they will become increasingly dependent on their supercomputers, just as in the 1970s, when businesses began to lean heavily on their mainframes. As prices continue to fall and usability improves, even smaller companies, especially those that are part of big-company supply chains, will be able to tap the increased efficiency made possible by supercomputers.

TAMING THE REBEL COMPUTER

There will be some culture shock, even outright resistance, perhaps, before supercomputers come to be seen as natural extensions of Fortune 500 computer rooms and data centers. Your average chief information officer, coming from a background in business and traditional computing methods, will cock an unsure eye at technology born of the physical sciences and proven in mostly engineering applications.

I'll hedge a bit when it comes to financial firms. Wall Street quants understand the dollars generated by algorithms running on 131,000 processors at once. To the majority of IT managers, supercomputers are outliers, rebels that can still only be tamed by an experienced few. But eventually, supercomputers, one of the fastest-growing segments of the computer hardware market, (9.4% compound annual growth rate), will be embraced by the traditional CIO community.

These large computers will even find their way to the smallest manufacturers. The Ohio Supercomputer Center, for example, offers the state's welders and machinists supercomputer time via the Internet to model welding techniques and prototypes for their larger clients. After year one of a pilot with a group of 200 welders, the results have been encouraging, cutting from weeks to hours the amount of time needed to run a typical simulation of welding a new material.

A NEW REVOLUTION

Some hurdles remain, notably the dearth of commercial software. Yet, 10 years since Deep Blue's historic chess game, signs bode well for the widespread commercial adoption of supercomputers.

The last silicon revolution saw PCs invade the office and automation take over the factory floor. During that period—from 1977 to 2001—overall U.S. manufacturing output almost doubled.

While productivity for the U.S. economy as a whole increased by 53%, manufacturing productivity rose 109%. Now we're on the cusp of another, more potent revolution powered by a highly developed form of silicon—the Age of the Supercomputer.

Let the games begin.

Dave Turek is the executive in charge of supercomputer development for IBM. He was a key designer of the IBM line of supercomputers which included Deep Blue, world chess champion, retired. Turek is a member of the U.S. Council on Competitiveness High Performance Computing Advisory Committee.

Age of the super brains. Focus Supercomputers are creating a stir as lightning-fast virtual analyzers, but the race for the next great breakthrough in IT has already begun â€“ with quantum computers. Find out how this development is changing society and business. 01 Supercomputers Concentrated computing power for simulations of the future. Page 12. 02 Quantum computers On the way to the next breakthrough in IT. Supercomputing is moving rapidly from its traditional place in academic and government research to tackle a growing list of commercial applications in business. Of the top 500 supercomputers in existence in 1997, only 161 were employed in businesses. Today, nearly half of the 500 speediestâ€”246 systemsâ€”reside in commercial enterprises. Two essential economic forces are propelling that movement. The first is the ever-rising expense of building and testing physical prototype designs for everything from diapers to cars. The second is the rapidly decreasing cost of perfectly modeling and testing a ...Â Now we're on the cusp of another, more potent revolution powered by a highly developed form of siliconâ€”the Age of the Supercomputer. Let the games begin. History of Supercomputing. Exploring todayâ€™s technology for tomorrowâ€™s possibilities.Â One of the first supercomputers to use integrated circuits. Performs calculations as fast as 160 megaFLOPS. Notable for its â€œCâ€ shape, which improved the efficiency of the hardware. The Cray X-MP. Released in 1982. Added 4 processors and greater memory bandwidth. Performs calculations as fast as 800 megaFLOPS. The Cray-2. Released in 1985. Cray, considered the â€œfather of supercomputing,â€ had left his post at business computing giant Sperry-Rand to join the newly formed Control Data Corporation so that he can focus on developing scientific computers. The title of worldâ€™s fastest computer was held at the time by the IBM 7030 â€œStretch,â€ one of the first to use transistors instead of vacuum tubes.Â Needless to say, Cray and his designs ruled the early era of the supercomputer. But he wasnâ€™t the only one advancing the field. The early â€™80s also saw the emergence of massively parallel computers, powered by thousands of processors all working in tandem to smash though performance barriers. A supercomputer is a computer with a high level of performance as compared to a general-purpose computer. The performance of a supercomputer is commonly measured in floating-point operations per second (FLOPS) instead of million instructions per second (MIPS). Since 2017, there are supercomputers which can perform over 1017 FLOPS (a hundred quadrillion FLOPS, 100 petaFLOPS or 100 PFLOPS). Since November 2017, all of the world's fastest 500 supercomputers run Linux-based operating systems. Additional